

# Mark Ibrahim

mark.ibhm@gmail.com • (973) 459-8429 • [markibrahim.me](http://markibrahim.me)  
in [marksibrahim](https://marksibrahim) • [github.com/marksibrahim](https://github.com/marksibrahim)

---

## TECHNICAL

• Python (scikit-learn, numpy, object-oriented, PyTest), Git, TensorFlow, Linux bash, SQL, JavaScript, LaTeX, VIM  
*familiar: Java, AWS, Neo4j, Elasticsearch, D3.js, Spark. Interests: Deep Learning Interpretability, Knowledge Graphs*

## EXPERIENCE

New York, NY

**Senior Machine Learning Engineer**, Center for Machine Learning at Capital One Sep 2016 – Present

- Leading Explainable AI team to build tools and research for explaining black-box deep learning models
  - Building open-source Python library to generate global explanations for neural network predictions
  - Published 2 interpretability research papers (NeurIPS workshop 2018 and ACM AAI 2019)
- Engineering a real-time [notification](#) system for predicting mistaken charges on 10 million transactions per day
  - Architecting Lambda microservices in Java and Python to serve machine learning predictions in production
- Developing recurrent neural network (RNN) + LDA customer archetype model in partnership with Columbia U.

**Data Engineering Fellow (2016) & Technical Advisor**, Insight Data Science May 2016 – Present

- Developed a graph-based knowledge search engine ([knowledgesearch.us](http://knowledgesearch.us)) powered by Wikipedia
  - Distributed parsing of all 5 million articles using Spark on Amazon Web Services (AWS)
- Designed a D3.js user interface powered by a graph database (Neo4j), Elasticsearch, and Python (Flask)

**Freelance Software Engineer**, Condé Nast Oct 2014 – Aug 2015

- Created an Applescript and Python [app](#) to tag and shorten Facebook/Twitter posts reaching 4 million followers
- Built a Google Calendar Extension to sync production sheet across team of writers/editors at ArchDigest

**Quantitative Portfolio Risk Analyst**, UBS Jun 2012 – Aug 2014

- Applied unsupervised machine learning (PCA) to identify \$570k in uncaptured sensitivity to 0.01% move in rates
- Automated daily 2½ hour manual risk calculation for \$658 million trading portfolio in Python

## SELECT RESEARCH

“Towards Explainable Deep Learning for Credit Lending”—*C Modarres, M Ibrahim, M Louie, J Paisley. NeurIPS 2018.*

“Global Explanations of Neural Networks: Mapping the Landscape of Predictions”—*M Ibrahim et al. AAI 2019.*

“Mixed Membership Recurrent Neural Networks”—*G Fazel, M Ibrahim, C Modarres, K Wu, J Paisley. Preprint 2019.*

“Connecting Every Bit of Knowledge: [Wikipedia’s First Link Network](#)”—*M Ibrahim et al. J. Computational Science 2017*

**Select Talks:** AAI Spotlight Talk (2019), NeurIPS FEAP Spotlight Workshop Talk (2018), NYC Python Meetup (2018), Tom Tom Machine Learning Conf (2018), [Data Driven](#) at George Washington U. (2017), NYC Data Wranglers (2017).

## COMMUNITY

**Reviewer** for academic journal *IEEE Transactions on Network Science and Engineering*, 2017-2018.

**Mentor** for Columbia U. Data Science Masters Capstone (2018). **Co-organizer** Vermont Python User Group (2016)

## EDUCATION

**M.S. Applied Mathematics**, University of Vermont (2016) Burlington, VT

Course Instructor: [Calculus I](#) (72 students) and [Calculus II](#) (38 students)

**Honors B.A. Mathematics, Magna Cum Laude**, Hamilton College (2012) Clinton, NY

[19th Gold Scholar](#) for student of “highest standards.” *Phi Sigma Iota*: highest honor for foreign languages